arXiv:1603.05496v1 [math.PR] 17 Mar 2016

Gain Function Approximation in the Feedback Particle Filter

Amirhossein Taghvaei, Prashant G. Mehta

***Abstract*— This paper is concerned with numerical algorithms for gain function approximation in the feedback particle filter. The exact gain function is the solution of a Poisson equation involving a probability-weighted Laplacian. The problem is to approximate this solution using *only* particles sampled from the probability distribution. Two algorithms are presented: a Galerkin algorithm and a kernel-based algorithm. Both the algorithms are adapted to the samples and do not require approximation of the probability distribution as an intermediate step. The paper contains error analysis for the algorithms as well as some comparative numerical results for a non-Gaussian distribution. These algorithms are also applied and illustrated for a simple nonlinear filtering example.**

I. Introduction

This paper is concerned with algorithms for numerically approximating the solution of a certain linear partial differential equation (pde) that arises in the problem of nonlinear filtering. In continuous time, the filtering problem pertains to the following stochastic differential equations (sdes):

$$\mathrm{d}X_t = a(X_t)\,\mathrm{d}t + \mathrm{d}B_t, \tag{1a}$$

$$\mathrm{d}Z_t = h(X_t)\,\mathrm{d}t + \mathrm{d}W_t, \tag{1b}$$

where $X_t \in \mathbb{R}^d$ is the (hidden) state at time t, $Z_t \in \mathbb{R}$ is the observation, and $\{B_t\}$, $\{W_t\}$ are two mutually independent standard Wiener processes taking values in $\mathbb{R}^d$ and $\mathbb{R}$, respectively. The mappings $a(\cdot):\mathbb{R}^d \to \mathbb{R}^d$ and $h(\cdot):\mathbb{R}^d \to \mathbb{R}$ are C^1 functions. Unless noted otherwise, all probability distributions are assumed to be absolutely continuous with respect to the Lebesgue measure, and therefore will be identified with their densities. The choice of observation being scalar-valued ($Z_t \in \mathbb{R}$) is made for notational ease.

The objective of the filtering problem is to estimate the posterior distribution of X_t given the time history of observations (filtration) $\mathcal{Z}_t := \sigma(Z_s : 0 \le s \le t)$. The density of the posterior distribution is denoted by p^*, so that for any measurable set $A \subset \mathbb{R}^d$,

$$\int_{x\in A} p^*(x,t)\,\mathrm{d}x = \mathsf{P}[X_t \in A \mid \mathcal{Z}_t].$$

The filter is infinite-dimensional since it defines the evolution, in the space of probability measures, of $\{p^*(\cdot,t) : t \ge 0\}$. If $a(\cdot)$, $h(\cdot)$ are linear functions, the solution is given by the finite-dimensional Kalman-Bucy filter. The article [3] surveys numerical methods to approximate the nonlinear filter. One approach described in this survey is particle filtering.

The particle filter is a simulation-based algorithm to approximate the filtering task [14]. The key step is the construction of N stochastic processes $\{X_t^i : 1 \le i \le N\}$: The value $X_t^i \in \mathbb{R}^d$ is the state for the i-th particle at time t. For each time t, the empirical distribution formed by the particle population is used to approximate the posterior distribution. Recall that this is defined for any measurable set $A \subset \mathbb{R}^d$ by,

$$p^{(N)}(A,t) = \frac{1}{N}\sum_{i=1}^{N} 1[X_t^i \in A].$$

A common approach in particle filtering is called *sequential importance sampling*, where particles are generated according to their importance weight at every time step [1], [14].

In our earlier papers [17], [16], [18], an alternative feedback control-based approach to the construction of a particle filter was introduced. The resulting particle filter, referred to as the feedback particle filter (FPF), is a controlled system. The dynamics of the i-th particle have the following gain feedback form,

$$\mathrm{d}X_t^i = a(X_t^i)\,\mathrm{d}t + \mathrm{d}B_t^i + \mathsf{K}_t(X_t^i) \circ \big(\mathrm{d}Z_t - \frac{h(X_t^i)+\hat{h}_t}{2}\,\mathrm{d}t\big), \tag{2}$$

where $\{B_t^i\}$ are mutually independent standard Wiener processes and $\hat{h}_t := \mathsf{E}[h(X_t^i)|\mathcal{Z}_t]$. The initial condition X_0^i is drawn from the initial density $p^*(x,0)$ of X_0 independent of $\{B_t^i\}$. Both $\{B_t^i\}$ and $\{X_0^i\}$ are also assumed to be independent of X_t, Z_t. The $\circ$ indicates that the sde is expressed in its Stratonovich form.

The gain function K_t is obtained by solving a weighted Poisson equation: For each fixed time t, the function ϕ is the solution to a Poisson equation,

$$\text{BVP} \qquad \begin{aligned} &\nabla\cdot(p(x,t)\nabla\phi(x,t)) = -(h(x)-\hat{h})p(x,t), \\ &\int \phi(x,t)p(x,t)\,\mathrm{d}x = 0 \qquad \text{(zero-mean)}, \end{aligned} \tag{3}$$

for all $x \in \mathbb{R}^d$ where ∇ and $\nabla\cdot$ denote the gradient and the divergence operators, respectively, and p denotes the conditional density of X_t^i given $\mathcal{Z}_t$.

In terms of the solution ϕ, the gain function is given by,

$$\mathsf{K}_t(x) = \nabla\phi(x,t)\,.$$

Note that the gain function K_t is vector-valued (with dimension $d\times 1$) and it needs to be obtained for each fixed time t. For the linear Gaussian case, the gain function is the Kalman gain. For the general nonlinear non-Gaussian problem, the FPF (2) is exact, given an exact gain function and an exact initialization $p(\cdot,0) = p^*(\cdot,0)$. Consequently, if the initial

Financial support from the NSF CMMI grants 1334987 and 1462773 is gratefully acknowledged.

A. Taghvaei and P. G. Mehta are with the Coordinated Science Laboratory and the Department of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign (UIUC). taghvae2@illinois.edu; mehtapg@illinois.edu;

conditions $\{X_0^i\}_{i=1}^N$ are drawn from the initial density $p^*(\cdot,0)$ of X_0, then, as $N\to\infty$, the empirical distribution of the particle system approximates the posterior density $p^*(\cdot,t)$ for each t.

A numerical implementation of the FPF (2) requires a numerical approximation of the gain function K_t and the mean $\hat{h}_t$ at each time-step. The mean is approximated empirically, $\hat{h}_t \approx \frac{1}{N}\sum_{i=1}^N h(X_t^i) =: \hat{h}_t^{(N)}$. The gain function approximation - the focus of this paper - is a challenging problem because of two reasons: i) Apart from the linear Gaussian case, there are no known closed-form solutions of (3); ii) The density $p(x,t)$ is not explicitly known. At each time-step, one only has samples X_t^i. These are assumed to be i.i.d sampled from p. Apart from the FPF algorithm, solution of the Poisson equation is also central to a number of other algorithms for nonlinear filtering [5], [15].

In our prior work, we have obtained results on existence, uniqueness and regularity of the solution to the Poisson equation, based on certain weak formulation of the Poisson equation. The weak formulation led to a Galerkin numerical algorithm. The main limitation of the Galerkin is that the algorithm requires a pre-defined set of basis functions - which scales poorly with the dimension d of the state. The Galerkin algorithm can also exhibit certain numerical issues related to the Gibb's phenomena. This can in turn lead to numerical instabilities in simulating the FPF.

The contributions of this paper are as follows: We present a new basis-free kernel-based algorithm for approximating the solution of the gain function. The key step is to construct a Markov matrix on a certain graph defined on the space of particles $\{X_t^i\}_{i=1}^N$. The value of the function ϕ for the particles, $\phi(X_t^i)$, is then approximated by solving a fixed-point problem involving the Markov matrix. The fixed-point problem is shown to be a contraction and the method of successive approximation applies to numerically obtain the solution.

We present results on error analysis for both the Galerkin and the kernel-based method. These results are illustrated with the aid of an example involving a multi-modal distribution. Finally, the two methods are compared for a filtering problem with a non-Gaussian distribution.

In the remainder of this paper, we express the linear operator in (3) as a weighted Laplacian $\Delta_\rho\phi := \frac{1}{\rho}\nabla\cdot(\rho\nabla\phi)$ where additional assumptions on the density ρ appear in the main body of the paper. In recent years, this operator and the associated Markov semigroup have received considerable attention with several applications including spectral clustering, dimensionality reduction, supervised learning etc [4], [10]. For a mathematical treatment, see the monographs [8], [2]. Related specifically to control theory, there are important connections with stochastic stability of Markov operators [7], [13].

The outline of this paper is as follows: The mathematical preliminaries appear in Sec. II. The Galerkin and the kernel-based algorithms for the gain function approximation appear in Sec. III and Sec. IV, respectively. The nonlinear filtering example appears in Sec. V.

II. Mathematical Preliminaries

The Poisson equation (3) is expressed as,

$$\text{BVP}\qquad \begin{aligned} -\Delta_\rho\phi &= h,\\ \int \phi\rho\,\mathrm{d}x &= 0 \qquad \text{(zero-mean)}, \end{aligned} \tag{4}$$

where ρ is a probability density on $\mathbb{R}^d$, $\Delta_\rho\phi := \frac{1}{\rho}\nabla\cdot(\rho\nabla\phi)$ and, without loss of generality, it is assumed $\hat{h} = \int h\rho\,\mathrm{d}x = 0$.

Problem statement: Approximate the solution $\phi(X^i)$ and $\nabla\phi(X^i)$ given N independent samples X^i drawn from ρ. The density ρ is not explicitly known.

For the problem to be well-posed requires definition of the function spaces and additional assumptions on ρ and h enumerated next: Throughout this paper, μ is an absolutely continuous probability measure on $\mathbb{R}^n$ with associated density ρ. $L^2(\mathbb{R}^d,\mu)$ is the Hilbert space of square integrable functions on $\mathbb{R}^d$ equipped with the inner-product,

$$<\phi,\psi> := \int \phi(x)\psi(x)\,\mathrm{d}\mu(x).$$

The associated norm is denoted as $\|\phi\|_2^2 := <\phi,\phi>$. The space $H^1(\mathbb{R}^d,\mu)$ is the space of square integrable functions whose derivative (defined in the weak sense) is in $L^2(\mathbb{R}^d,\mu)$. We use L^2 and H^1 to denote $L^2(\mathbb{R}^d,\mu)$ and $H^1(\mathbb{R}^d,\mu)$, respectively. For the zero-mean solution of interest, we additionally define the co-dimension 1 subspace $L_0^2 := \{\phi\in L^2; \int\phi\,\mathrm{d}\mu = 0\}$ and $H_0^1 := \{\phi\in H^1; \int\phi\,\mathrm{d}\mu = 0\}$. L^∞ is used to denote the space of functions that are bounded a.e. (Lebesgue) and the sup-norm of a function $\phi\in L^\infty$ is denoted as $\|\phi\|_\infty$.

The following assumptions are made throughout the paper:

(i) **Assumption A1:** The probability density function ρ is of the form $\rho(x) = e^{-(x-\mu)^T\Sigma^{-1}(x-\mu)-V(x)}$ where $\mu\in\mathbb{R}^d$, Σ is a positive-definite matrix and $V\in C^2$ with $V\in L^\infty$, and its derivatives $DV, D^2V\in L^\infty$.

(ii) **Assumption A2:** The function $h\in L^2$ and $\int h\,\mathrm{d}\mu = 0$.

Under assumption A1, the density ρ admits a spectral gap (or Poincaré inequality)[2], i.e $\exists\,\lambda_1>0$ such that,

$$\int\phi^2\rho\,\mathrm{d}x \le \frac{1}{\lambda_1}\int|\nabla\phi|^2\rho\,\mathrm{d}x,\quad \forall\,\phi\in H_0^1. \tag{5}$$

Furthermore, the spectrum is known to be discrete with an ordered sequence of eigenvalues $0=\lambda_0<\lambda_1\le\lambda_2\le\cdots$ and associated eigenfunctions $\{e_n\}$ that form a complete orthonormal basis of L^2 [Corollary 4.10.9 in [2]]. The trivial eigenvalue $\lambda_0 = 1$ with associated eigenfunction $e_0 = 1$. On the subspace of zero-mean functions, the spectral decomposition yields: For $\phi\in L_0^2$,

$$-\Delta_\rho\phi = \sum_{m=1}^\infty \lambda_m <e_m,\phi> e_m. \tag{6}$$

The spectral gap condition (5) implies that $\lambda_1>0$. Consequently, the semigroup

$$e^{t\Delta_\rho}\phi := \sum_{m=1}^\infty e^{-t\lambda_m}<e_m,\phi>e_m \tag{7}$$

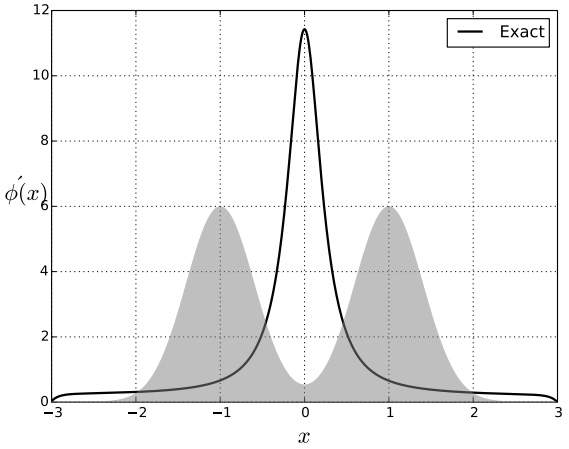

Fig. 1. The exact solution to the Poisson equation using the formula (8). The density ρ is the sum of two Gaussians $N(-1,\sigma^2)$ and $N(+1,\sigma^2)$, and $h(x) = x$. The density is depicted as the shaded curve in the background.

is a strict contraction on the subspace L_0^2. It is also easy to see that μ is an invariant measure and $\int e^{t\Delta_\rho}\phi(x)\,\mathrm{d}\mu(x) = \int \phi(x)\,\mathrm{d}\mu(x) = 0$ for all $\phi \in L_0^2$.

Example 1: If the density ρ is Gaussian with mean $\mu \in \mathbb{R}^d$ and a positive-definite covariance matrix Σ, the spectral gap constant $(1/\lambda_1)$ equals the largest eigenvalue of Σ. The eigenfunctions are the Hermite functions. Given a linear function $h(x) = H \cdot x$, the unique solution of the BVP (4) is given by,

$$\phi(x) = (\Sigma H)\cdot(x-\mu),$$

where $\cdot$ denotes the vector dot product in $\mathbb{R}^d$. In this case, $\mathsf{K} = \nabla\phi = \Sigma H$ is the Kalman gain. ■

Example 2: In the scalar case (where $d = 1$), the Poisson equation is:

$$-\frac{1}{\rho(x)}\frac{\mathrm{d}}{\mathrm{d}x}(\rho(x)\frac{\mathrm{d}\phi}{\mathrm{d}x}(x)) = h.$$

Integrating twice yields the solution explicitly,

$$\begin{aligned}\frac{\mathrm{d}\phi}{\mathrm{d}x}(x) &= -\frac{1}{\rho(x)}\int_{-\infty}^{x}\rho(z)h(z)\,\mathrm{d}z \\ \phi(x) &= -\int_{-\infty}^{x}\frac{\mathrm{d}y}{\rho(y)}\int_{-\infty}^{y}\rho(z)h(z)\,\mathrm{d}z\end{aligned} \tag{8}$$

For the particular choice of ρ as the sum of two Gaussians $N(-1,\sigma^2)$ and $N(+1,\sigma^2)$ with $\sigma = 0.4$ and $h(x) = x$, the solution obtained using (8) is depicted in Fig. 1. Since $\frac{\mathrm{d}h}{\mathrm{d}x} > 0$, the positivity of the gain function $\frac{\mathrm{d}\phi}{\mathrm{d}x}(x)$ follows from the maximum principle for elliptic pdes [6]. ■

III. Galerkin for Gain Function Approximation

Weak formulation: A function $\phi \in H_0^1$ is said to be a weak solution of Poisson's equation (4) if

$$<\nabla\phi, \nabla\psi> = <h,\psi>, \quad \forall\ \psi \in H_0^1. \tag{9}$$

It is shown in [12] that, under Assumptions (A1)-(A2), there exists a unique weak solution of the Poisson equation.

The Galerkin approximation involves solving (9) in a finite-dimensional subspace $S \subset H_0^1(\mathbb{R}^d;\rho)$. The solution ϕ is approximated as,

$$\phi^{(M)}(x) = \sum_{m=1}^{M} c_m \psi_m(x), \tag{10}$$

where $\{\psi_m(x)\}_{m=1}^M$ are a *given* set of basis functions.

The finite-dimensional approximation of (9) is to choose constants $\{c_m\}_{m=1}^M$ such that

$$<\nabla\phi^{(M)}, \nabla\psi> = <h,\psi>, \ \forall\ \psi \in S, \tag{11}$$

where $S := \mathrm{span}\{\psi_1,\cdots,\psi_M\} \subset H_0^1$.

Denoting $[A]_{ml} = <\nabla\psi_l\cdot\nabla\psi_m>$, $b_m = <h,\psi_m>$, and $c = (c_1,c_2,\cdots,c_M)^T$, the finite-dimensional approximation (11) is expressed as a linear matrix equation:

$$Ac = b. \tag{12}$$

In a numerical implementation, the matrix A and vector b are approximated as,

$$[A]_{ml} = <\nabla\psi_l\cdot\nabla\psi_m> \approx \frac{1}{N}\sum_{i=1}^{N}\nabla\psi_l(X^i)\cdot\nabla\psi_m(X^i) =: [A]_{ml}^{(N)}, \tag{13}$$

$$b_m = <h,\psi_m> \approx \frac{1}{N}\sum_{i=1}^{N} h(X^i)\psi_m(X^i) =: b_m^{(N)}. \tag{14}$$

The resulting solution of the matrix equation (12), with $A = A^{(N)}$ and $b = b^{(N)}$, is denoted as $c^{(N)}$,

$$A^{(N)}c^{(N)} = b^{(N)}. \tag{15}$$

Using (10), we obtain the particle-based approximation of the solution:

$$\phi^{(M,N)}(x) = \sum_{m=1}^{M} c_m^{(N)}\psi_m(x). \tag{16}$$

In terms of this solution, the gain function is obtained as,

$$\nabla\phi^{M,N}(x) = \sum_{m=1}^{M} c_m^{(N)}\nabla\psi_m(x).$$

Convergence analysis: The following is a summary of the approximations in the Galerkin algorithm:

$$\begin{aligned}\text{Exact:}\qquad & <\nabla\phi,\nabla\psi> = <h,\psi>, \quad \forall\,\psi\in H_0^1 \\ \text{Galerkin approx:}\qquad & <\nabla\phi^{(M)},\nabla\psi> = <h,\psi>,\ \forall\,\psi\in S\subset H_0^1\end{aligned}$$

Empirical approximation:

$$\frac{1}{N}\sum_{i=1}^{N}\nabla\phi^{(M,N)}(X^i)\cdot\nabla\psi(X^i) = \frac{1}{N}\sum_{i=1}^{N}h(X^i)\psi(X^i),\ \ \forall\,\psi\in S$$

We are interested in the error analysis of these approximations as a function of both M and N. Note that the approximation error $\phi^{(M)} - \phi^{(M,N)}$ is random because X^i are sampled randomly from the probability distribution μ. The following Proposition provides error bounds for the

case where the basis functions are the eigenfunctions of the Laplacian. The proof appears in the Appendix A.

Proposition 1: Consider the empirical Galerkin approximation of the Poisson equation (9) on the space $S := \text{span}\{e_1, e_2, \cdots, e_M\}$ of the first M eigenfunctions. Fix $M < \infty$. Then the solution for the matrix equation (15) exists with probability approaching 1 as $N \to \infty$. And there is a sequence of random variables $\{\varepsilon_N\}$ such that

$$\|\phi^{(M,N)} - \phi\|_2 \le \frac{1}{\lambda_M^2}\|h\|_2^2 + \varepsilon_N, \tag{17}$$

where $\varepsilon_N \to 0$ as $N \to \infty$ a.s. ■

Remark 1: In practice, the eigenfunctions of the Laplacian are not known. The basis functions are typically picked from the polynomial family, e.g., the Hermite functions. In this case, the bounds will provide qualitative assessment of the error provided the eigenfunctions associated with the first M eigenvalues are 'approximately' in S. Quantitatively, the additional error may be bounded in terms of the projection error between the eigenfunction and the projection onto S. ■

Example 3: We next revisit the bimodal distribution first introduced in the Example 2. Fig. 2 depicts the empirical Galerkin approximation, with $N = 200$ samples, of the gain function. The basis functions are from the polynomial family - $\{x, x^2, \cdots, x^M\}$ - and the figure depicts the gain functions with $M = 1, 3, 5$ modes. The $(M = 1)$ case is referred to as the *constant-gain approximation* where,

$$\frac{\mathrm{d}\phi}{\mathrm{d}x}(x) = \int h(x)\, x \, \mathrm{d}\mu(x) = (\text{const.})$$

For the linear Gaussian case, the (const.) is the Kalman gain. The plots included in the Figure demonstrate the Gibb's phenomena as more and more modes are included in the Galerkin. Particularly concerning is the change in the sign of the gain which leads to negative values of gain for some particles contradicting the positivity property of the gain described in Example 2. As discussed in the filtering example in Sec. V, the negative values of the gain can cause numerical issues and lead to erroneous results for the filter. ■

IV. Semigroup Approximation of the Gain

Semigroup formulation: The semigroup formula (7) is used to construct the solution of the Poisson equation by solving the following fixed-point equation for *any* fixed positive value of t:

$$\phi = e^{t\Delta_\rho}\phi + \int_0^t e^{s\Delta_\rho} h \, \mathrm{d}s. \tag{18}$$

A unique solution exists because $e^{t\Delta_\rho}$ is a contraction on L_0^2.

The approximation proposed in this section involves approximating the semigroup as a perturbed integral operator, for small positive values of $t = \varepsilon$. The following approximation of the semigroup appears in [4], [10]:

$$T^{(\varepsilon)}\phi(x) := \frac{\int k^{(\varepsilon)}(x,y)\phi(y)\,\mathrm{d}\mu(y)}{\int k^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(y)}, \tag{19}$$

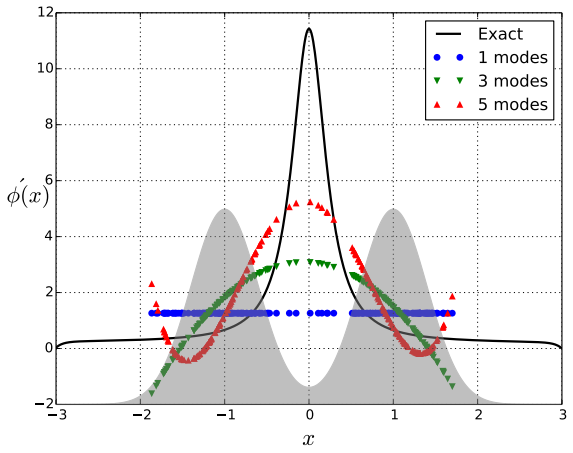

Fig. 2. Comparison of the exact solution and its empirical Galerkin approximation with $M = 1, 3, 5$ modes and $N = 200$ particles. The density is depicted as the shaded curve in the background.

where $k^{(\varepsilon)}(x,y) := \frac{g^{(\varepsilon)}(x,y)}{\sqrt{\int g^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(y)}\sqrt{\int g^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(x)}}$ and $g^{(\varepsilon)}(x,y) := \frac{1}{(4\pi\varepsilon)^{\frac{d}{2}}}\exp\left(-\frac{|x-y|^2}{4\varepsilon}\right)$ is the Gaussian kernel in $\mathbb{R}^d$.

In terms of the perturbed integral operator, the fixed-point equation (18) becomes,

$$\phi^{(\varepsilon)} = T^{(\varepsilon)}\phi^{(\varepsilon)} + \int_0^\varepsilon T^{(s)} h \, \mathrm{d}s. \tag{20}$$

The superscript ε is used to distinguish the approximate (ε-dependent) solution from the exact solution ϕ. As shown in the Appendix B, $T^{(\varepsilon)}$ has an ergodic invariant measure $\mu^{(\varepsilon)}$ which approximates μ as $\varepsilon \downarrow 0$. For any fixed positive ε, we are interested in solutions that are zero-mean with respect to this measure. The existence-uniqueness result for this solution is described next; the proof appears in the Appendix B.

Proposition 2: Consider the fixed-point problem (20) with the perturbed operator $T^{(\varepsilon)}$ defined according to (19). Fix $\varepsilon > 0$. Then there exists a unique solution $\phi^{(\varepsilon)}$ such that $\int \phi^{(\varepsilon)}\,\mathrm{d}\mu^{(\varepsilon)} = 0$. ■

In a numerical implementation, the solution ϕ^ε is approximated directly for the particles:

$$\Phi^{(\varepsilon,N)} := (\phi^{(\varepsilon)}(X^1), \phi^{(\varepsilon)}(X^2), \cdots, \phi^{(\varepsilon)}(X^N)).$$

The integral operator $T^{(\varepsilon)}$ is approximated as a $N \times N$ Markov matrix whose (i,j) entry is obtained empirically as,

$$T_{i,j}^{(\varepsilon,N)} = \frac{k^{(\varepsilon,N)}(X^i, X^j)}{\sum_{l=1}^N k^{(\varepsilon,N)}(X^i, X^l)}, \tag{21}$$

where $k^{(\varepsilon,N)}(x,y) = \frac{g^\varepsilon(x,y)}{\sqrt{\frac{1}{N}\sum_{l=1}^N g^\varepsilon(x,X^l)}\sqrt{\frac{1}{N}\sum_{l=1}^N g^\varepsilon(y,X^l)}}$.

The resulting finite-dimensional fixed-point equation is given by,

$$\Phi^{(\varepsilon,N)} = T^{(\varepsilon,N)}\Phi^{(\varepsilon,N)} + \int_0^\varepsilon T^{(s,N)} H^{(N)} \, \mathrm{d}s, \tag{22}$$

where $\Phi^{(\varepsilon,N)} = (\Phi_1^{(\varepsilon,N)}, \Phi_2^{(\varepsilon,N)}, \cdots, \Phi_N^{(\varepsilon,N)}) \in \mathbb{R}^N$ is the vector-valued solution, $H^{(N)} = (h(X^1), h(X^2), \cdots, h(X^N)) \in \mathbb{R}^N$, and $T^{(\varepsilon,N)}$ and $T^{(s,N)}$ are $N \times N$ matrices defined according to (21). The existence-uniqueness result for the zero-mean solution of the finite-dimensional equation (22) is described next; its proof is given in the Appendix C. The zero-mean property is the finite-dimensional counterpart of the zero-mean condition $\int \phi \, \mathrm{d}\mu = 0$ for the original problem and $\int \phi \, \mathrm{d}\mu^\varepsilon = 0$ for the perturbed problem.

Proposition 3: Consider the fixed-point problem (22) with the matrix $T^{(\varepsilon,N)}$ defined according to (21). Then, with probability 1, there exists a unique zero-mean solution $\Phi^{(\varepsilon,N)} \in \mathbb{R}^d$. ■

Once $\Phi^{(\varepsilon,N)}$ is available, it is straightforward to extend it to the entire domain. For $x \in \mathbb{R}^d$,

$$\phi^{(\varepsilon,N)}(x) := \frac{\sum_{i=1}^N k^{(\varepsilon,N)}(x, X^i)\Phi_i^{(\varepsilon,N)}}{\sum_{i=1}^N k^{(\varepsilon,N)}(x, X^i)} + \int_0^\varepsilon T^{(s,N)} h(x) \, \mathrm{d}s. \tag{23}$$

By construction $\phi^{(\varepsilon,N)}(X^i) = \Phi_i^{(\varepsilon,N)}$ for $i = 1, \cdots, N$. The extension is not necessary for filtering because one only needs to solve for $\nabla\phi$ at X^i. The formula for this is,

$$\frac{\partial \phi}{\partial x_l}(X^i) = \sum_{j=1}^N \left[T_{ij}^{(\varepsilon,N)} \Phi_j^{(\varepsilon,N)} \left(X_l^j - \sum_{k=1}^N T_{ik}^{(\varepsilon,N)} X_l^k \right) \right]$$

where $X_l^i \in \mathbb{R}$ is the l-th component of $X^i \in \mathbb{R}^d$.

The algorithm is summarized in Algorithm 1. For numerical purposes we use $\varepsilon H^N \approx \int_0^\varepsilon T^{(s,N)} H^{(N)} \, \mathrm{d}s$. Also, the fixed-point problem (22) is conveniently solved using the method of successive approximation. In filtering, the initial guess is readily available from the solution at the previous time-step.

Algorithm 1 Kernel-based gain function approximation

Input: $\{X^i\}_{i=1}^N$, $H := \{h(X^i)\}_{i=1}^N$, ε
Output: $\Phi := \{\phi(X^i)\}_{i=1}^N$, $\{\nabla\phi(X^i)\}_{i=1}^N$

Calculate $g_{ij} := \exp(-|X^i - X^j|^2/4\varepsilon)$ for $i, j = 1$ to N.

Calculate $k_{ij} := \frac{g_{ij}}{\sqrt{\sum_l g_{il}}\sqrt{\sum_l g_{jl}}}$ for $i, j = 1$ to N.

Calculate $T_{ij} := \frac{k_{ij}}{\sum_l k_{il}}$ for $i, j = 1$ to N.

Solve $\Phi = T\Phi + \varepsilon H$ for Φ (Successive approximation).

Calculate $\frac{\partial\phi}{\partial x_l}(X^i) = \sum_{j=1}^N \left[T_{ij}\Phi_j \left(X_l^j - \sum_{k=1}^N T_{ik} X_l^k \right) \right]$ for $l = 1$ to d.

Convergence: The following is a summary of the approximations with the kernel-based method:

$$\text{Exact:} \qquad \phi = e^{\varepsilon\Delta_\rho}\phi + \int_0^\varepsilon e^{s\Delta_\rho} h \, \mathrm{d}s$$

$$\text{Kernel approx:} \qquad \phi^{(\varepsilon)} = T^{(\varepsilon)}\phi^{(\varepsilon)} + \int_0^\varepsilon T^{(s)} h \, \mathrm{d}s$$

$$\text{Empirical approx:} \qquad \phi^{(\varepsilon,N)} = T^{(\varepsilon,N)}\phi^{(\varepsilon,N)} + \int_0^\varepsilon T^{(s,N)} h \, \mathrm{d}s$$

We break the convergence analysis into two steps. The first step involves convergence of $\phi^{(\varepsilon,N)}$ to $\phi^{(\varepsilon)}$ as $N \to \infty$. The second step involves convergence of $\phi^{(\varepsilon)}$ to ϕ as $\varepsilon \to 0$. The following theorem states the convergence result for the first step; a sketch of the proof appears in Appendix D.

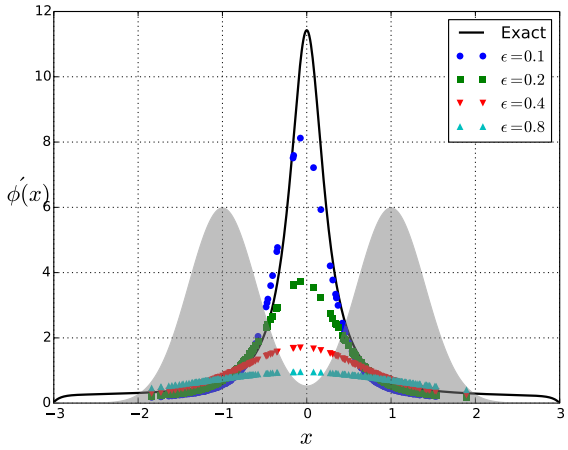

Fig. 3. Comparison of the exact solution and its kernel-based approximation with $\varepsilon = 0.1, 0.2, 0.4, 0.8$ and $N = 200$ particles. The density is depicted as the shaded curve in the background.

Theorem 1: Consider the empirical kernel approximation of the fixed-point equation (18). Fix $\varepsilon > 0$. Then,

(i) There exists a unique (zero-mean) solution $\phi^{(\varepsilon)}$ for the perturbed fixed-point equation (20).

(ii) For any finite N, a unique (zero-mean) solution $\Phi^{(\varepsilon,N)}$ for (22) exists with probability 1.

For a compact set $\Omega \subset \mathbb{R}^d$,

$$\lim_{N\to\infty} \sup_{x\in\Omega} |\phi^{(\varepsilon,N)}(x) - \phi^{(\varepsilon)}(x)| = 0, \quad \text{a.s.,} \tag{24}$$

where $\phi^{(\varepsilon,N)}$ is the extension of the vector-valued solution $\Phi^{(\varepsilon,N)}$ (see (23)). ■

The convergence analysis for step 2, as $\varepsilon \to 0$, is the subject of ongoing work. In this regard, it is shown in [9] that for compactly supported functions $f \in C^3$,

$$f(x) - T^{(\varepsilon)} f(x) = \varepsilon \Delta_\rho f(x) + O(\varepsilon^2).$$

Example 4: Consider once more the bimodal distribution introduced in 2. Figure 3 depicts the kernel-based approximation of the gain function with $N = 200$ particles and range of ε. The kernel-based avoids the Gibbs phenomena observed with the Galerkin (see Fig. 2). Notably, the gain function is positive for any choice of ε. ■

V. Numerics

In this section, we consider a filtering problem associated with the bimodal distribution introduced in Example 2. The filtering model is given by,

$$\mathrm{d}X_t = 0,$$
$$\mathrm{d}Z_t = X_t \, \mathrm{d}t + \sigma_w \, \mathrm{d}W_t,$$

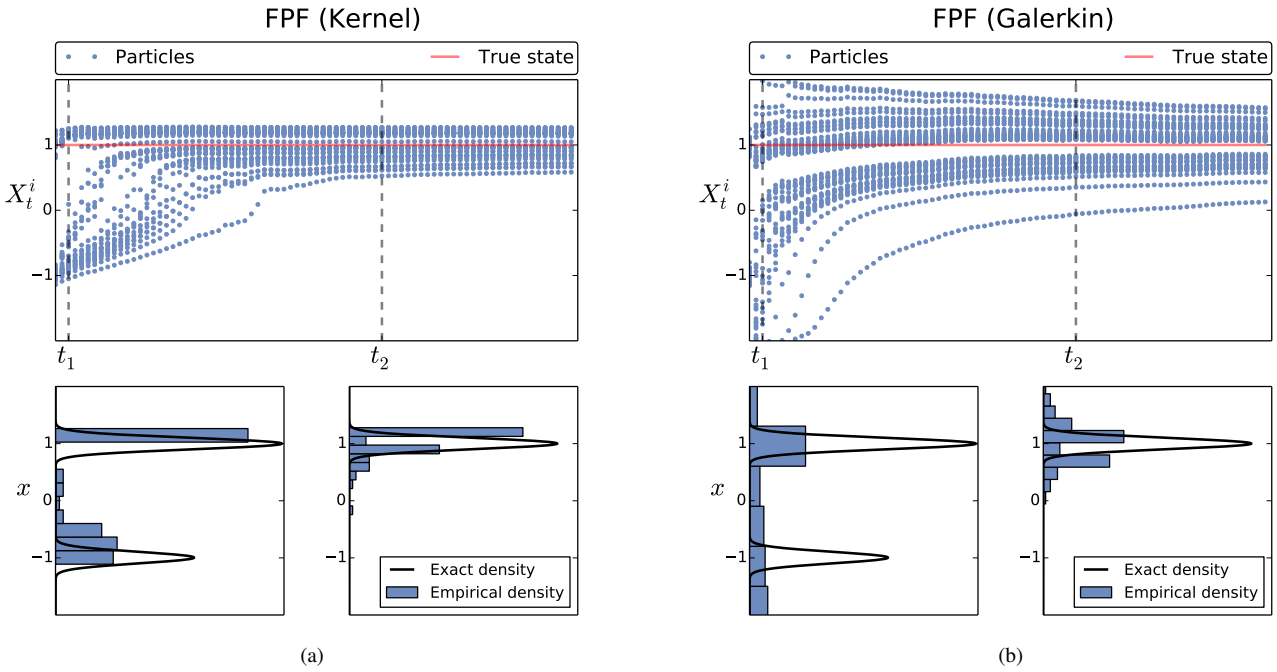

Fig. 4. Comparison of simulation results: Trajectory of particles and posterior distribution with (a) Kernel-based approximation of the gain function, (b) Galerkin approximation of the gain function.

where $X_t \in \mathbb{R}$, $Z_t \in \mathbb{R}$, W_t is a standard Wiener process, the initial condition X_0 is sampled from the bimodal distribution comprising of two Gaussians, $N(-1,\sigma^2)$ and $N(+1,\sigma^2)$, and without loss of generality, $Z_0 = 0$. As in Example 2, the observation function $h(x) = x$ is linear. The static case is considered because the posterior is given explicitly:

$$p^*(x,t) \doteq (\text{const.}) \exp\left(\frac{1}{2\sigma_W^2} h(x) Z_t - \frac{1}{4\sigma_W^2} |h(x)|^2 t \right) p_0^*(x). \tag{25}$$

The following filtering algorithms are implemented for comparison:

1) Kalman filter;
2) Feedback particle filter with the Galerkin approximation where $S = \text{span}\{x, x^2, x^3, x^4, x^5\}$;
3) Feedback particle filter with the kernel approximation.

The performance metrics are as follows:

1) Filter mean $\hat{X}_t$;
2) Conditional probability $\mathsf{P}\left[|X_t - X_0| < \frac{1}{2}|\mathcal{Z}_t\right]$.

The simulation parameters are as follows: The true initial state $X_0 = 1$. The measurement noise parameter $\sigma_W = 0.3$. The simulation is carried out over a finite time-horizon $t \in [0,T]$ with $T = 0.8$ and a fixed discretized time-step $\Delta t = 0.02$. All the particle filters use $N = 100$ particles and have the same initialization, where particles are drawn with equal probability from one of the two Gaussians, $N(-1,\sigma^2)$ or $N(+1,\sigma^2)$, where $\sigma = 0.1$. For the kernel approximation, we use $\varepsilon = 0.15$. The simulation parameters are tabulated in Table I. The Kalman filter is initialized with $\hat{X}_0 = 0$ and $\Sigma_0 = \text{Var}(X_0) = 1 + \sigma^2$. The latter corresponds to the variance of the prior. Figure 4 parts (a) and (b) depict the particle trajectories and the associated distributions obtained using the kernel approximation and the Galerkin approximation, respectively. The kernel-based approximation provides for a better approximation of the exact posterior. At time t_1 during the initial transients, some of the particles with the Galerkin approximation show a divergence. This is a numerical issue due to the Gibb's phenomena that leads to erroneous negative value of the gain (see the discussion in Examples 2 and 3).

Figure 5 depicts a comparison of the simulation results for the two metrics. For the particle filters, these are computed empirically.

For applications of FPF with kernel-based approximation of the gain function for attitude estimation problem see the companion paper [19].

References

[1] A. Bain and D. Crisan. *Fundamentals of stochastic filtering*, volume 3. Springer, 2009.

[2] D. Bakry, I. Gentil, and M. Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013.

[3] A. Budhiraja, L. Chen, and C. Lee. A survey of numerical methods for nonlinear filtering problems. *Physica D: Nonlinear Phenomena*, 230(1):27–36, 2007.

[4] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.

[5] F. Daum, J. Huang, and A. Noushin. Exact particle flow for nonlinear filters. In *SPIE Defense, Security, and Sensing*, pages 769704–769704, 2010.

[6] L. C. Evans. *Partial Differential Equations*. American Mathematical Society, Providence, RI, 1998.

[7] P. W. Glynn and S. P. Meyn. A Liapunov bound for solutions of the Poisson equation. *The Annals of Probability*, pages 916–931, 1996.

[8] A. Grigoryan. Heat kernels on weighted manifolds and applications. *Cont. Math*, 398:93–191, 2006.

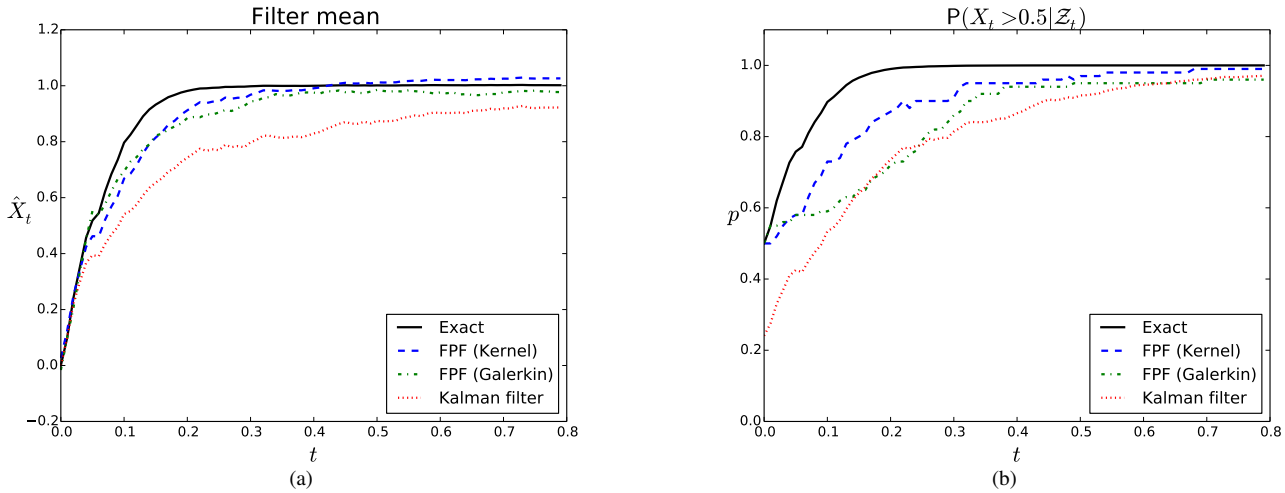

Fig. 5. Comparison of simulation results: (a) Conditional mean of the state X_t, (b) Conditional probability of X_t larger than $\frac{1}{2}$.

[9] M. Hein, J. Audibert, and U. Von Luxburg. From graphs to manifolds–weak and strong pointwise consistency of graph Laplacians. In *Learning theory*, pages 470–485. Springer, 2005.

[10] M. Hein, J. Audibert, and U. Von Luxburg. Graph Laplacians and their convergence on random neighborhood graphs. *arXiv preprint math/0608522*, 2006.

[11] V. Hutson, J. Pym, and M. Cloud. *Applications of functional analysis and operator theory*, volume 200. Elsevier, 2005.

[12] R. S. Laugesen, P. G. Mehta, S. P. Meyn, and M. Raginsky. Poisson's equation in nonlinear filtering. *SIAM Journal on Control and Optimization*, 53(1):501–525, 2015.

[13] S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.

[14] A. Smith, A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo methods in practice*. Springer Science & Business Media, 2013.

[15] T. Yang, H. Blom, and P. G. Mehta. The continuous-discrete time feedback particle filter. In *American Control Conference (ACC), 2014*, pages 648–653. IEEE, 2014.

[16] T. Yang, P. G. Mehta, and S. P Meyn. Feedback particle filter with mean-field coupling. In *Decision and Control and European Control Conference (CDC-ECC), 50th IEEE Conference on*, pages 7909–7916, 2011.

[17] T. Yang, P. G. Mehta, and S. P. Meyn. A mean-field control-oriented approach to particle filtering. In *American Control Conference (ACC), 2011*, pages 2037–2043. IEEE, 2011.

[18] T. Yang, P. G. Mehta, and S. P. Meyn. Feedback particle filter. *IEEE Trans. Automatic Control*, 58(10):2465–2480, October 2013.

[19] C. Zhang, A. Taghvaei, and P. G. Mehta. Attitude estimation with feedback particle filter. Available online: http://mehta.mechse.illinois.edu/downloads/conferences/2016_CDC_Attitude_Estimation_FPF.pdf.

APPENDIX

A. Proof of Proposition 1

Using the spectral representation (6), because $h \in L^2$,

$$\phi = -\Delta_\rho^{-1} h = \sum_{m=1}^{\infty} \frac{1}{\lambda_m} < e_m, h > e_m.$$

With basis functions as eigenfunctions,

$$\phi^{(M)} = \sum_{m=1}^{M} \frac{1}{\lambda_m} < e_m, h > e_m.$$

Therefore,

$$\|\phi - \phi^{(M)}\|_2^2 = \sum_{m=M+1}^{\infty} \frac{1}{\lambda_m^2} |< e_m, h >|^2 \le \frac{1}{\lambda_M^2} \|h\|_2^2.$$

Next, by the triangle inequality,

$$\|\phi - \phi^{(M,N)}\|_2 \le \|\phi - \phi^{(M)}\|_2 + \|\phi^{(M)} - \phi^{(M,N)}\|_2.$$

We show $\|\phi^{(M)} - \phi^{(M,N)}\|_2 \to 0$ a.s. as $N \to \infty$.

Using the formulae (10) and (16) with basis-functions ψ_m as eigenfunctions e_m,

$$\|\phi^{(M)} - \phi^{(M,N)}\|_2 = |c - c^N|,$$

where $|\cdot|$ denotes the Euclidean norm in $\mathbb{R}^M$. The vectors c and $c^{(N)}$ solve the matrix equations (see (12) and (15)),

$$Ac = b,$$
$$A^{(N)} c^{(N)} = b^{(N)},$$

where $A^{(N)} \xrightarrow{\text{a.s.}} A$, $b^{(N)} \xrightarrow{\text{a.s.}} b$ by the strong law of large numbers. Consequently, because $A = \mathrm{diag}(\lambda_1, \ldots, \lambda_N)$ is invertible, $c^{(N)} \to c$ a.s. as $N \to \infty$.

B. Proof of Proposition 2

Denote $n^{(\varepsilon)}(x) := \int k^{(\varepsilon)}(x,y)\, \mathrm{d}\mu(y)$, and re-write the operator $T^{(\varepsilon)}$ as,

$$T^{(\varepsilon)} f(x) = \int \frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x) n^{(\varepsilon)}(y)} f(y)\, \mathrm{d}\mu^{(\varepsilon)}(y),$$

where $\mathrm{d}\mu^{(\varepsilon)}(x) := n^{(\varepsilon)}(x)\, \mathrm{d}\mu(x)$.

Denote $L^2(\mu^{(\varepsilon)})$ as the space of square integrable functions with respect to $\mu^{(\varepsilon)}$ and as before $L_0^2(\mu^{(\varepsilon)}) := \left\{ \phi \in L^2(\mu^{(\varepsilon)}) \middle| \int \phi\, \mathrm{d}\mu^{(\varepsilon)} = 0 \right\}$ is the co-dimension 1 subspace of mean-zero functions in $L^2(\mu^{(\varepsilon)})$. The technical part of proving the Proposition is to show that the operator $T^{(\varepsilon)}$ is a strict contraction on the subspace.

Lemma 1: Suppose ρ, the density of the probability measure μ, satisfies Assumption A1. Then,

(i) $\mu^{(\varepsilon)}$ is a finite measure.

(ii) For sufficiently small values of ε, the operator $T^{(\varepsilon)} : L^2(\mu^{(\varepsilon)}) \to L^2(\mu^{(\varepsilon)})$ is a compact Markov operator with an invariant measure $\mu^{(\varepsilon)}$.

(iii) $T^{(\varepsilon)} : L_0^2(\mu^{(\varepsilon)}) \to L_0^2(\mu^{(\varepsilon)})$ is a strict contraction.

Proof:

(i) WLOG assume $\mu = 0$ in the Assumption A1. For notational ease, denote

$$\rho^{(\varepsilon)}(x) := \int g^{(\varepsilon)}(x,y)\rho(y)\,\mathrm{d}y,$$

where recall that $\rho^{(\varepsilon)}(x)$ is used to define the denominator of the kernel. Then

$$c_1 \exp(-\tfrac{1}{2}x^T Q_1^{-1} x \le \rho^{(\varepsilon)}(x) \le c_2 \exp(-\tfrac{1}{2}x^T Q_1^{-1} x),$$

where $Q_1^{-1} := (\Sigma + 2\varepsilon I)^{-1}$ and $c_1 = (2\pi)^{-\frac{d}{2}}|Q_1|^{-\frac{1}{2}} e^{-\|V\|_\infty}$ and $c_2 = (2\pi)^{-\frac{d}{2}}|Q_1|^{-\frac{1}{2}} e^{\|V\|_\infty}$ are positive constants that depend on $\|V\|_\infty$. Therefore,

$$\begin{aligned}\mu^{(\varepsilon)}(\mathbb{R}^d) &= \int\int k^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(x)\,\mathrm{d}\mu(y)\\ &= \int\int \frac{g^{(\varepsilon)}(x,y)}{\sqrt{\rho^{(\varepsilon)}(x)}\sqrt{\rho^{(\varepsilon)}(y)}}\rho(x)\rho(y)\,\mathrm{d}x\,\mathrm{d}y\\ &\le \frac{1}{c_1}\int\int g^{(\varepsilon)}(x,y)\, e^{-\frac{1}{2}x^T Q_2^{-1}x} e^{-\frac{1}{2}y^T Q_2^{-1}y}\,\mathrm{d}x\mathrm{d}y,\end{aligned}$$

which is bounded because $Q_2^{-1} := \Sigma^{-1} - \frac{1}{2}Q_1^{-1} > 0$. This proves that $\mu^{(\varepsilon)}$ is a finite measure.

(ii) The integral operator T^ε is a Markov operator because the kernel $k^{(\varepsilon)}(x,y) > 0$ and

$$T^{(\varepsilon)}1 = \int \frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x)n^{(\varepsilon)}(y)}\,\mathrm{d}\mu^{(\varepsilon)}(y) = \int \frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x)}\,\mathrm{d}\mu(y) = 1.$$

T^ε is compact because [Theorem 7.2.7 in [11]],

$$\int\int \left|\frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x)n^{(\varepsilon)}(y)}\right|^2 \mathrm{d}\mu^{(\varepsilon)}(x)\,\mathrm{d}\mu^{(\varepsilon)}(y) < \infty.$$

The inequality holds because the integrand can be bounded by a Gaussian:

$$\exp(-\frac{|x-y|^2}{2\varepsilon} - \tfrac{1}{2}x^T Q_3^{-1} x^T - \tfrac{1}{2}y^T Q_3 y)$$

where $Q_3^{-1} := \Sigma^{-1} - \frac{1}{2}Q_1^{-1} - (Q_2 + 2\varepsilon I)^{-1}$ is of order $O(\varepsilon)$.

Finally, the measure μ^ε is an invariant measure because for all functions f,

$$\int T^{(\varepsilon)}f(x)\ \mathrm{d}\mu^{(\varepsilon)}(x) = \int f(y)\ \mathrm{d}\mu^{(\varepsilon)}(y).$$

(iii) Since $\mu^{(\varepsilon)}$ is an invariant measure, $T^{(\varepsilon)} : L_0^2(\mu^{(\varepsilon)}) \to L_0^2(\mu^{(\varepsilon)})$. Next, for all $f, g \in L_0^2(\mu^{(\varepsilon)})$,

$$\begin{aligned}&\tfrac{1}{2}\int f(x)^2\,\mathrm{d}\mu^{(\varepsilon)}(x) + \tfrac{1}{2}\int g(y)^2\,\mathrm{d}\mu^{(\varepsilon)}(y)\\ &- \int\int \frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x)n^{(\varepsilon)}(y)} f(x)g(y)\,\mathrm{d}\mu^{(\varepsilon)}(x)\,\mathrm{d}\mu^{(\varepsilon)}(y)\\ &= \int\int \frac{k^{(\varepsilon)}(x,y)}{n^{(\varepsilon)}(x)n^{(\varepsilon)}(y)}(f(x)-g(y))^2\,\mathrm{d}\mu^{(\varepsilon)}(x)\,\mathrm{d}\mu^{(\varepsilon)}(y)\\ &\ge 0,\end{aligned}$$

where, because the kernel is everywhere positive, the equality holds iff

$$f(x) = g(y) = (\text{const.}) \quad \mu^\varepsilon - \text{a.e.}$$

Therefore, substituting $g = T^{(\varepsilon)}f$ in the inequality above,

$$\|T^{(\varepsilon)}f\|^2 \le \frac{1}{2}\|f\|^2 + \frac{1}{2}\|T^{(\varepsilon)}f\|^2 \Rightarrow \|T^{(\varepsilon)}f\| \le \|f\|,$$

where the L^2 norms here are with respect to the invariant measure μ^ε. Now, for $f \in L_0^2(\mu^{(\varepsilon)})$, $\int f\,\mathrm{d}\mu^{(\varepsilon)} = 0$, and thus the equality holds iff

$$f(x) = (\text{const.}) = 0.$$

Therefore, $T^{(\varepsilon)}$ is strictly contractive on $L_0^2(\mu^{(\varepsilon)})$. ∎

The proof of the Prop. 2 now follows because $T^{(\varepsilon)}$ is a contraction on $L_0^2(\mu^{(\varepsilon)})$. Note also that $\mu^{(\varepsilon)}(A) \to \mu(A)$ for any measurable set, as $\varepsilon \downarrow 0$.

C. Proof of Proposition 3

By construction, $T^{(\varepsilon,N)}$ is a $N \times N$ stochastic matrix whose entries are all positive with probability 1. The result follows.

D. Proof sketch for Theorem 1

Parts (i) and (ii) have already been proved as part of the Prop. 2 and Prop. 3, respectively. The convergence result (24) leans on approximation theory for integral operators. In [Theorem 7.6.6 in [11]], it is shown that if

(i) $T^{(\varepsilon)}$ is compact and $T^{(\varepsilon,N)}$ is collectively compact;

(ii) $\|T^{(\varepsilon)}\| < 1$;

(iii) $T^{(\varepsilon,N)}$ converges to $T^{(\varepsilon)}$ pointwise, i.e

$$\lim_{N\to\infty} \|T^{(\varepsilon,N)}f - T^{(\varepsilon)}f\|_\infty = 0 \quad \text{a.s,} \quad \forall f$$

Then for N large enough, $(I - T^{(\varepsilon,N)})^{-1}$ is bounded and

$$\lim_{N\to\infty} \|(I - T^{(\varepsilon,N)})^{-1}h - (I - T^{(\varepsilon)})^{-1}h\|_\infty = 0, \quad \text{a.s} \quad \forall h.$$

In order to prove the convergence result, we consider the Banach space $C_0(\Omega) := \{f \in C(\Omega);\ \int f\,\mathrm{d}\mu(x) = 0\}$ equipped with the $\|\cdot\|_\infty$ norm. We consider $T^{(\varepsilon)}$ and $T^{(\varepsilon,N)}$ as linear operators on $C_0(\Omega)$. Note that $T^{(\varepsilon,N)}$ here corresponds to the extension of the vector-valued solution to $\mathbb{R}^d$ using (23).

The proof involves verification of each of the three requirements stated above:

(i) Collective compactness follows because $k^{(\varepsilon)}$ is continuous [Theorem 7.2.6 [11]].

(ii) The norm condition holds because

$$\begin{aligned}|T^{(\varepsilon)}f(x)| &\le \frac{\int k^{(\varepsilon)}(x,y)|f(y)|\,\mathrm{d}\mu(y)}{\int k^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(y)}\\ &\le \frac{\int k^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(y)}{\int k^{(\varepsilon)}(x,y)\,\mathrm{d}\mu(y)}\|f\|_\infty \le \|f\|_\infty\end{aligned}$$

where the equality holds only when f is constant. For $f \in C_0(\Omega)$, this constant can only be 0.

(iii) Pointwise convergence follows from using the LLN. For a fixed continuous function f, LLN implies convergence for every fixed $x \in \mathbb{R}^d$. On a compact set Ω, pointwise convergence also implies uniform convergence.

∎